

An Advanced On-Board Processor and Adaptive Scanning Controller for the Next-Generation Precipitation Radar

Andrew C. Berkun, Mark A. Fischman, and Eastwood Im

Jet Propulsion Laboratory, California Institute of Technology, Radar Science and Engineering Section

M/S 300-241, 4800 Oak Grove Dr, Pasadena, CA 91109-8099, USA

Tel: (818) 354-0742, Fax: (818) 393-6943, e-mail: aberkun@ampersand.jpl.nasa.gov

Abstract—The Next-Generation Precipitation Radar (PR-2) prototyped by NASA/JPL will depend heavily on high-performance digital processing to collect meaningful rain echo data. Using field-programmable gate arrays (FPGAs), we have developed for the PR-2 a pulse-compression processor and adaptive timing controller that will enable full on-board processing capabilities in a 14 and 36 GHz spaceborne radar. This paper describes some of the new technologies for the on-board processor, including a 40×10^9 op/s bit-serial filter attaining -60 dB range sidelobe performance, and an adaptive scanning control and timing unit (CTU) which yields a 7-fold increase in the radar's dwell time over areas of precipitation.

I. INTRODUCTION

Global three-dimensional rainfall data acquired by a spaceborne precipitation radar provide crucial information about the hydrological cycle. The imaging of liquid water mass in three-dimensions (in altitude and horizontally) is also important for understanding the global climate system because such images reveal how latent heat is transferred through the atmosphere. The Ku-band (14 GHz) precipitation radar onboard the Tropical Rainfall Measuring Mission (TRMM) satellite, launched in 1997, was the first Earth-orbiting sensor to retrieve 3-D rainfall data over tropical and mid-latitude regions [1]. With the success of TRMM, a new class of spaceborne precipitation radars are being developed for follow-on missions with improved sensitivity, spatial resolution, and swath coverage.

NASA/JPL is currently prototyping a Next-Generation Precipitation Radar (PR-2) that offers the advanced capabilities [2] needed to succeed TRMM. The PR-2 project involves the parallel development of several breakthrough technologies [3], including: 1) a large aperture membrane reflector antenna, 2) compact, solid-state T/R modules for an electronically steerable array at 14 and 36 GHz, and 3) on-board digital data processing and timing control, implemented in field-programmable gate arrays (FPGAs). The on-board processor, which is the focus of this paper, must be designed to meet the challenging requirements for measuring precipitation from space. Sea surface returns can be as much as 55 dB stronger than the rain return power to become a major source of clutter [4]. Also, for conventional scanning radars, there is a classic tradeoff between swath width capabilities and the available dwell time over each footprint [5]. In the following sections, we describe the new features of our prototype data processor and control and timing unit (CTU) that solve these system design challenges.

II. PULSE COMPRESSION DATA PROCESSOR

The PR-2 uses a chirp radar technique, which has advantages over short pulse weather radar systems like TRMM that typically sample at a vertical range resolution of 250 m. With a chirp rain radar, the number of independent looks within the 250 m cell can be increased by modulating the chirp over several megahertz bandwidth. For example, the PR-2 uses a chirp bandwidth of $B = 4$ MHz—or equivalently a range resolution of $c/2B = 37.5$ m—to increase the number of samples within the range bin by a factor of $N = 6.7$. After these samples are averaged in power to reduce Rayleigh fading noise, the effective signal-to-noise ratio of the radar's rain echo measurement can be improved by as much as $\sqrt{N} = 4.1$ dB. A further benefit of chirp radar is that, because signal power is gained through the use of a long transmit pulse length ($\tau = 50$ μ s for the PR-2), the requirements for instantaneous transmit power are relatively low (~ 200 W). This moderate transmit power requirement allows simpler, solid-state transmitter hardware to be integrated into the phased array antenna feeds.

The cost of the chirp radar is that the pulse compression processing in the receiver section is numerically intensive. To attain sufficiently low range sidelobes of -60 dB for masking out sea surface clutter, a large matched filter bank is needed with 200 or more coefficients and with time-domain weighting of the chirp reference pulse [6]. Taking into account the complex I/Q arithmetic for the matched filter and the requirement for 4 separate receive channels (Ku and Ka-band, H and V-pol), the PR-2 must be able to process incoming data at 40 billion operations per second.

High speed, radiation-hardened FPGA chips with million gate densities have recently emerged which can support the high throughput requirements for the PR-2. In the last year, our team partner, Andra Consulting, Inc., has implemented our concept for a new pulse compression data processor (shown in Fig. 1) onto two FPGAs using a distributed arithmetic technique [7]. We have successfully tested the first version of the processor core aboard an airborne prototype of the PR-2, deployed during the 4th Convection and Moisture Experiment (CAMEX-4) in the summer of 2001. Fig. 2 shows the pulse-compressed and processed science data from a CAMEX-4 flight over Hurricane Humberto in September, 2001. As a next step toward a spaceborne PR-2 mission, the data processor is being re-designed to interface with a more sophisticated timing controller for satellite-based measurements, where a number of echoes-in-flight (EIF) must be tracked and averaged together over time.

The onboard processor is designed around a bit serial finite impulse response (FIR) filter implemented in a pair of Xilinx Virtex-1000 FPGAs. Data enters the FPGAs as four channels of 12-bit digitized data at a 20 MHz sampling rate. The 4 MHz chirp data is

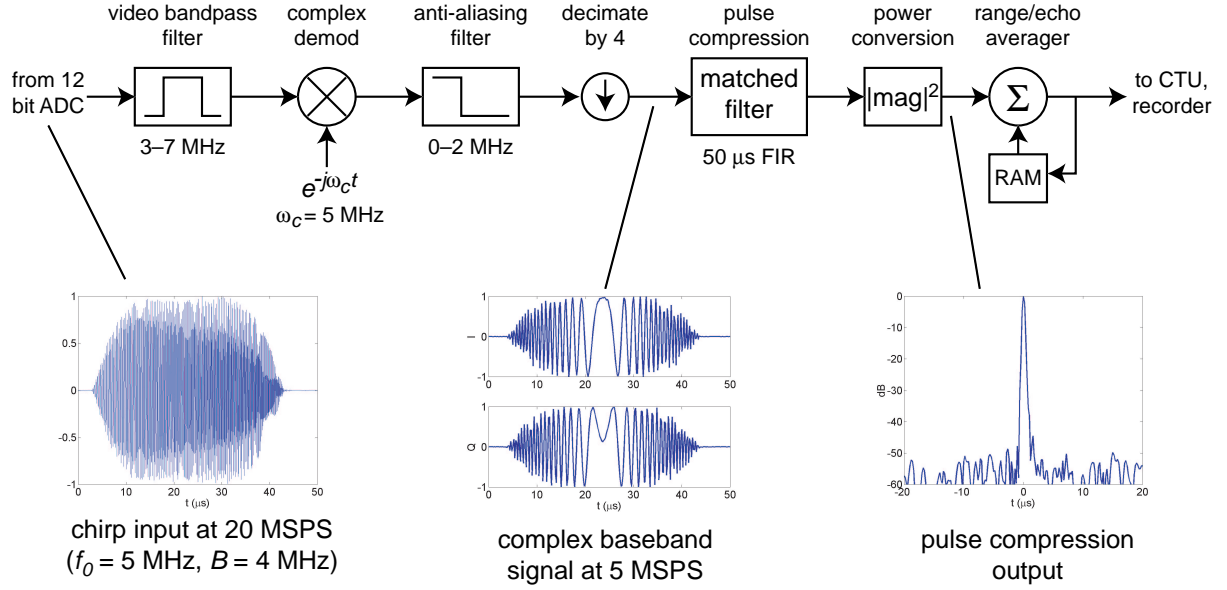


Figure 1: Block diagram of the digital data processor signal flow for the PR-2.

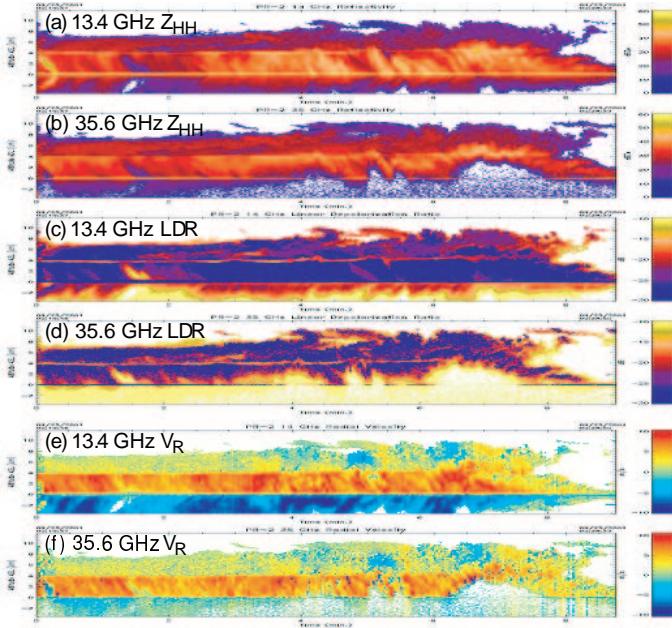


Figure 2: Field demonstration of the FPGA-based data processor aboard the airborne PR-2 prototype. The above graphs show dual-frequency (Ku and Ka-band), multi-parameter rain profiles taken over Hurricane Humberto on September 25, 2001, during the CAMEX-4 experiment. Graphs (a) and (b) display co-polarized, effective radar reflectivity data taken over a time scale of 8 minutes. (The bright line near the bottom of each graph represents the ocean surface return at 0 km altitude). Graphs (c) and (d) show the linear depolarization ratios, which accentuate scattering by ice particles at the melting layer. Graphs (e) and (f) track the vertical rainfall motion using doppler measurements of the echo returns.

centered around a 5 MHz offset video carrier. The remainder of the bandwidth is noise which has leaked through the front end filters.

The first stage of processing is a saturation detector followed by an anti-aliasing, symmetric FIR filter, with 64 taps and 16-bit coefficients, which removes the out-of-band noise and the alias centered at 15 MHz. The output of this filter is mixed with a 5 MHz sinusoid to shift the center of the signal spectrum down to 0 Hz. The data is then downsampled by a factor of 4 with no loss of information due to the anti-aliasing filter.

The 5 MHz complex baseband data is stored in a first-in/first-out memory so that pulses containing saturated samples can be eliminated. Data is then fed into a 256-tap, non-symmetric FIR filter which contains the complex conjugate of the expected return. This is where the bulk of the FPGA's processing resources are used, amounting to 20 billion real multiply and add operations per second. The output of this 256-tap filter is the compressed radar return. Complex data is converted to power and averaged over adjacent range bins and multiple pulses (Fig. 3). Averaging provides two orders of magnitude data compression. Further range averaging is applied for adaptive scanning, described in the next section, to identify the beams most likely to contain rain.

A number of DSP applications make this processing density possible. Bit-serial filtering makes best use of the speed capabilities of the FPGA. The FIR filter stage uses a lookup table to combine 4 stages of multiplication and addition into a single configurable logic block. The video filter, demodulator, and downsampler are actually four parallel filter stages which compute only the samples which are kept after decimation. Complex-valued filtering is implemented by running delay stages and multipliers at double speed and switching between real and imaginary components of each received waveform, saving half of the FPGA's real estate.

III. CONTROL AND TIMING UNIT

A unique feature of precipitation is that it is often sparsely distributed over the Earth's surface. Even for weather in tropical re-

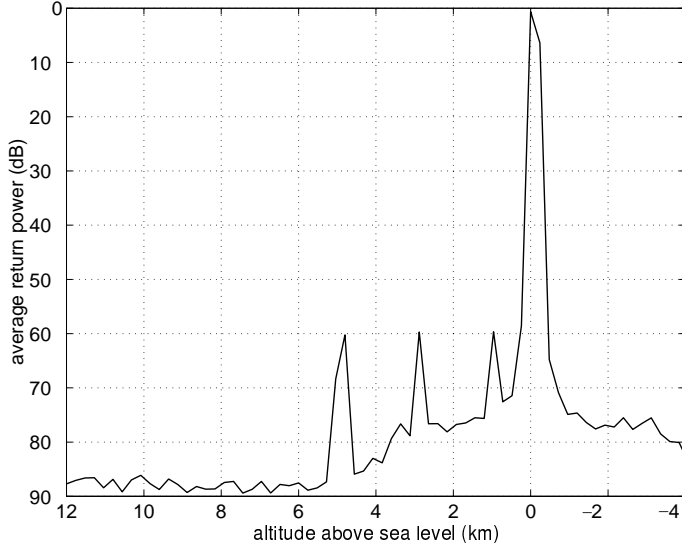


Figure 3: Hardware description language simulation of the spaceborne Data Processor FPGA for rain over the ocean. The pulse-compressed echo output shows that light rain targets at 1, 3, and 5 km altitudes are detectable even when these echoes are 60 dB below the sea surface return.

regions, which accounts for more than half of the total global rainfall, precipitation occurs over only 4% of the surface area [4]. This statistic can be exploited in the PR-2 by using an adaptive scanning technique [5]—that is, the radar can use its own preliminary “quick-scan” data to electronically steer the radar beam to only those areas which contain precipitation, and to ignore areas that are precipitation-free. With adaptive scanning, the dwell time for the rain target can be increased 7 times beyond that of a conventional cross-track scanning radar without sacrificing swath width.

We have developed a specialized control and timing unit with an adaptive scanning algorithm onto a Virtex-1000 FPGA. This CTU generates the transmit and receive timing for an entire 300 ms sweep cycle, consisting of: 1) a locator sweep that performs a rough measurement of echo return power over all cross-track beam locations (248 beams) and over an 8 km altitude range; 2) a bubble-sort algorithm which ranks the 248 beam locations from highest to lowest return powers; and 3) a high-resolution sweep, which takes additional radar looks of the top 24 ranked beams over a longer 12 km altitude range. The timing solution for the CTU must be generated on-the-fly with minimum dead-time between EIFs, while also avoiding collisions between transmit and receive echoes.

Because of the critical timing constraints of the adaptive scan, we chose to implement the CTU with a network of custom-made state machines in the FPGA rather than with a microprocessor core. The advantage to using state machines is that they can respond immediately and in-parallel to a large number of timing interrupts. Fig. 4 shows a diagram of the key logic modules developed for the CTU. The sweep engine module includes an EIF counter array which keeps track of the location of up to 32 radar echoes and provides feedback to guarantee that transmit and receive pulses are interlaced. A bubble-sort machine takes echo return values from the locator sweep and executes an $N^2/4$ sorting algorithm in the block RAM of the FPGA. An antenna driver module then sends beam steering data from an SRAM lookup table to 2,300 phase-shifters in the PR-2’s

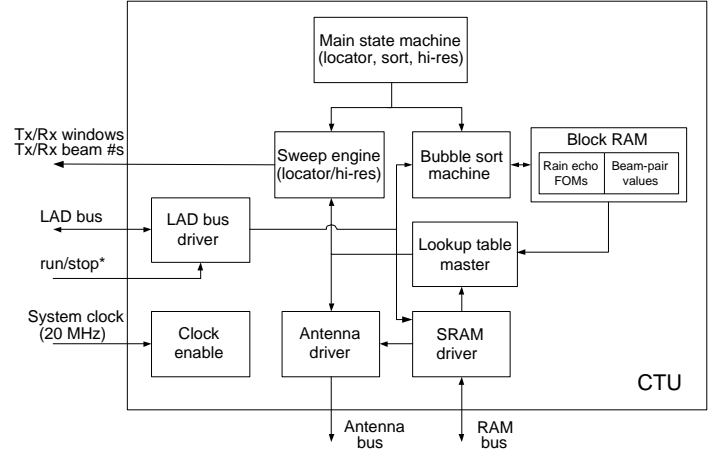


Figure 4: Verilog module design for the Control and Timing Unit.

antenna array to define the present locations of transmit and receive beams. All of the digital circuitry for the CTU has been designed at a register transfer level using the Verilog hardware description language. Currently this design has a total equivalent gate count of 73,000 gates. To guarantee that setup and hold timing requirements are met, we implemented the CTU logic as a synchronous design (with all registers tied to a common 20 MHz clock) and ran it through a static timing analysis.

The timing performance of the CTU was tested using Verilog simulation software. For a simplified test case with 16 cross-track beams scanning $\pm 37^\circ$ from nadir and with a shortened sweep cycle of 16.1 ms, 77% of the total sweep time was dedicated to either transmitting a pulse or receiving an echo. This chip-level simulation underestimates the true timing efficiency due to the abbreviated sweep cycle; a greater amount of time is spent either building up EIFs at the beginning of the sweep or receiving the remaining EIFs near the end of the sweep. From our calculations of the PR-2’s pulse repetition sequence over a full 300 ms cycle and all 248 beam locations, the estimated timing efficiency will actually range from 87–94%, depending on how the rain is distributed. To test the speed of the bubble-sort processor, several random rain profiles were generated in software for a full 248 beam scan and were input to the CTU Verilog simulator. The sorting machine had an average run time of 2.0 ± 0.1 ms, which shows that ranking of rain locations can be completed within a small fraction (less than 1%) of the 300 ms sweep cycle.

IV. SPACE QUALIFIED DESIGN

The physical board design for the data processor and CTU provides the standard interfaces in a satellite instrument and uses equivalent S class parts. Key parts are the Virtex-1000, which has passed radiation testing and is being evaluated for S class, and the 15 MHz static RAM chips. Key interfaces include a 1553 spacecraft command bus, data output bus for a solid state recorder, telemetry, science ADC, the electrically steerable antenna bus currently under development, an arbitrary waveform generator for the chirp transmission, and a test port for ground-based operation. The board is also being mounted in a custom-made chassis, shown in Fig. 5, which is capable of removing the 20 watts of heat from each FPGA in a vacuum and stabilizing the board under 10 G’s of launch acceleration.

One design problem in radiation environments is that any memory

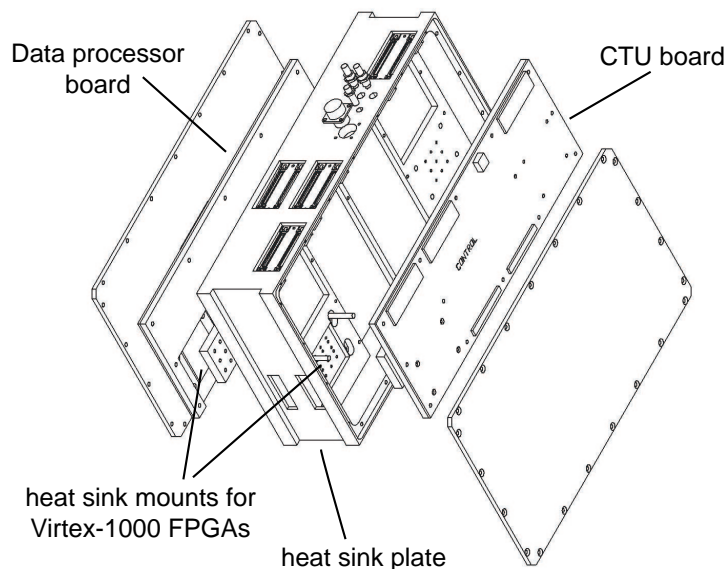


Figure 5: Mechanical chassis design for the spaceborne PR-2 processor/controller.

cell or register can be corrupted by a single event upset (SEU). Triple redundancy is commonly used to mitigate SEUs, but that approach was deemed too costly for the PR-2 application. The Virtex FPGA configuration is SRAM-based and is therefore susceptible to radiation upsets. The approach chosen is to use a master FPGA implemented as a one-time programmable, antifuse ACTEL part, which is immune to reconfiguration by radiation. The ACTEL part runs an error correction routine on a large SRAM bank using a Hamming code to correct single bit errors as they occur. Periodically this SRAM is used to reconfigure the Xilinx FPGAs and their command tables. In this way, an SEU cannot permanently disable the instrument. Multiple bit errors in the same word of the RAM between scrubbing cycles are extremely unlikely.

V. CONCLUSION

We have presented an overview of the Next-Generation Precipitation Radar's on-board digital electronics design, which includes a pulse compression processor with very low range sidelobe performance and a new adaptive scanning controller. The on-board processor/controller will serve as a key subsystem for a future satellite precipitation radar mission having dual-frequency (14 and 36 GHz) capabilities.

ACKNOWLEDGMENT

The research described in this paper was performed by the Jet Propulsion Laboratory, California Institute of Technology, under contract with the National Aeronautics and Space Administration.

REFERENCES

[1] T. Kozu, T. Kawanishi, H. Kuroiwa, M. Kojima, K. Oikawa, H. Kumagai, K. Okamoto, M. Okumura, H. Nakatsuka, and

K. Nishikawa, "Development of precipitation radar onboard the Tropical Rainfall Measuring Mission (TRMM) satellite," *IEEE Trans. Geosci. Remote Sensing*, vol. 39, pp. 102–116, Jan. 2001.

[2] E. Im, S. L. Durden, G. Sadowy, A. Berkun, J. Huang, M. Lou, B. C. Lopez, Y. Rahmat-Sammi, and S. Rengarajan, "System concept for the next-generation spaceborne precipitation radars," in *Proc. 2000 Aero. Conf.*, vol. 5, (Big Sky, MT), pp. 151–158, IEEE, Mar. 2000.

[3] G. Sadowy, A. Berkun, S. Durden, J. Huang, E. Im, B. Lopez, M. Lou, Y. Rahmat-Samii, F. Y. Liu, and S. Rengarajan, "Technologies for the next generation of spaceborne precipitation radars," in *Proc. 2001 Aero. Conf.*, vol. 4, pp. 4/1811–4/1823, IEEE, 2001.

[4] R. Meneghini and T. Kozu, *Spaceborne Weather Radar*. Norwood, MA: Artech House, 1990.

[5] D. Atlas, "Adaptively pointing spaceborne radar for precipitation measurements," *J. Appl. Meteor.*, vol. 21, pp. 429–431, Mar. 1982.

[6] A. Tanner, S. L. Durden, R. Denning, E. Im, F. K. Li, W. Ricketts, and W. Wilson, "Pulse compression with very low sidelobes in an airborne rain mapping radar," *IEEE Trans. Geosci. Remote Sensing*, vol. 32, pp. 211–213, Jan. 1994.

[7] R. Andraka and A. Berkun, "FPGAs make a radar signal processor on a chip a reality," in *Proc. 33rd Asilomar Conf. Signals, Syst. Comput.*, (Monterey, CA), IEEE, Oct. 1999.